Variable Geometry as a Multiplex Repeated Game: Overlapping Clubs, Hegemonic Bilateralism, and Cross-Issue Punishment

Mark Daley

February 1, 2026

Abstract

This note formalizes "variable geometry"—overlapping, issue-specific coalitions of middle powers—as a multiplex network of clubs interacting repeatedly in the presence of a hegemon that can offer targeted bilateral inducements. The baseline one-shot logic is a familiar pick-off problem: bilateral side-payments can unravel cooperation. The core result here is that multiplex overlap can endogenously create cross-issue punishment capacity: defection in one layer triggers exclusion or retaliation across a portfolio of layers, raising the minimum inducement required for deviation (a "divide-and-rule tax" on the hegemon). We specify a repeated game with portable reputations and excludable club benefits, derive subgame-perfect incentive constraints with proofs, characterize the hegemon's minimum bribes to induce defection, and show how threshold technologies (minimum viable club size) generate tipping: beyond a multiplex density threshold, buying defections becomes prohibitively expensive.

1 Motivation and Contribution

The premise of "variable geometry" is that, as universal multilateral institutions weaken, states form pragmatic, issue-specific coalitions that overlap across domains (security, procurement, standards, critical minerals, AI governance, etc.). The incentive problem is immediate: why would a middle power remain in a club if a hegemon offers an attractive bilateral deal?

A one-shot model makes clubs fragile. The interesting possibility is that overlap across clubs creates *cross-issue enforcement*: deviation in one arena damages standing in others. This paper provides a rigorous formalization of that mechanism and derives conditions under which it is incentive-compatible.

2 Model

2.1 Players, Issues, and Multiplex Architecture

There are N middle powers indexed by $i \in \mathcal{N}$ and one hegemon H. There are L issues (layers) indexed by $\ell \in \mathcal{L} = \{1, \ldots, L\}$.

Each layer ℓ is associated with a club membership set $S_\ell \subseteq \mathcal{N}$. Let $k_\ell = |S_\ell|$. Define the membership set of i:

$$\mathcal{L}_i \equiv \{\ell \in \mathcal{L} : i \in S_\ell\}, \qquad d_i \equiv |\mathcal{L}_i| \ \text{ (multiplex degree).}$$

Variable geometry means that the sets $\{S_\ell\}_{\ell \in \mathcal{L}}$ overlap but are not identical.

2.2 Timing Within a Period

Time is infinite, $t = 0, 1, 2, \ldots$. In each period t:

1. The hegemon simultaneously offers *contingent* transfers (bribes) $b_{i\ell}^t \geq 0$ to any $i \in S_\ell$, payable only if i chooses to defect in layer ℓ in that period.
2. In each layer ℓ, each member $i \in S_\ell$ chooses an action $a_{i\ell}^t \in \{C, D\}$.
3. Actions are publicly observed within each layer. A cross-layer information system may propagate deviations across layers (defined below).
4. Payoffs for period t are realized.

Players discount future payoffs with common factor $\delta \in (0, 1)$. (Nothing substantive hinges on a common δ, but it simplifies exposition.)

2.3 Stage Payoffs in a Layer

Fix a layer ℓ and a member $i \in S_\ell$. Define:

- $u_{i\ell}^C$: i's per-period payoff in layer ℓ when all members of S_ℓ cooperate and i is in good standing.
- $u_{i\ell}^P$: i's per-period payoff in layer ℓ when i is punished/excluded in that layer.
- $\pi_{i\ell} \geq 0$: i's *intrinsic* one-period deviation gain in ℓ from choosing D when others choose C, excluding the hegemonic transfer (saved compliance costs, opportunistic gains, etc.).

If i deviates unilaterally in ℓ (chooses D while others in S_ℓ choose C) and accepts transfer $b_{i\ell}^t$, then i's layer-ℓ payoff in period t is

$$u_{i\ell}^C + \pi_{i\ell} + b_{i\ell}^t.$$

Define the per-layer *excludable surplus*:

$$s_{i\ell} \equiv u_{i\ell}^C - u_{i\ell}^P > 0. \tag{1}$$

Total per-period payoff for i is separable across layers:

$$U_i^t \equiv \sum_{\ell \in \mathcal{L}_i} u_{i\ell}^t,$$

where $u_{i\ell}^t$ is determined by actions and punishment status in layer ℓ at time t.

2.4 Portable Reputation and Spillover Across Layers

Let $r_i^t \in \{G, B\}$ denote i's reputation at time t (good/bad). Initially, $r_i^0 = G$ for all i.

A deviation in any layer can "spill over" to other layers by making r_i bad and triggering punishment beyond the deviated layer.

Assumption 1 (Spillover and information-sharing)**.** *If i chooses D in any layer at time t, then:*

1. $r_i^{t+1} = B$ *(bad reputation).*
2. *For each other layer $\ell' \in \mathcal{L}_i \setminus \{\ell\}$, the deviation is recognized and punishment is applied in ℓ' with probability $\alpha \in [0, 1]$, independently across ℓ'. The deviated layer punishes with probability* 1.

The parameter α is an explicit measure of the "connective tissue" among clubs: transparency rules, intelligence sharing, institutional linkages, and political willingness to condition cooperation in one arena on behavior in another.

2.5 Credibility of Targeted Punishment

Punishment must be a best response for non-deviators. We impose a minimal (and realistic) structure: punishment is implemented by *exclusion from excludable club goods* (procurement access, standards recognition, stockpiles, preferential finance, etc.), which need not reduce payoffs of punishers.

Assumption 2 (Costless exclusion technology)**.** *For any layer ℓ and any member $i \in S_\ell$, there exists a punishment action profile by $S_\ell \setminus \{i\}$ that delivers $u_{i\ell}^P$ to i and delivers at least $u_{j\ell}^C$ to each $j \in S_\ell \setminus \{i\}$.*

Assumption 2 makes targeted punishment self-enforcing: punishers weakly prefer exclusion to continued inclusion of the deviator. This is exactly where "club goods" matter: if benefits are non-excludable, punishment is not credible.

3 Strategies: Multiplex Grim Trigger with Targeted Exclusion

Definition 1 (Multiplex grim trigger)**.** *Fix $\alpha \in [0,1]$. The following strategy profile is used by middle powers:*

- *If $r_i^t = G$, then in every layer $\ell \in \mathcal{L}_i$, play C and reject any transfer $b_{i\ell}^t$.*
- *If $r_i^t = B$, then in every layer $\ell \in \mathcal{L}_i$, accept the punished/excluded status and receive $u_{i\ell}^P$ (others implement exclusion as in Assumption 2).*
- *Upon observing a deviation by some i in any layer at time t, update $r_i^{t+1} = B$ and apply spillover punishments across layers as in Assumption 1.*

We will show when this strategy profile constitutes a subgame perfect equilibrium (SPE) in the repeated game among middle powers *given* a transfer offer process, and then characterize the hegemon's optimal inducements.

4 Equilibrium Analysis: Incentive Constraints and Proofs

Let

$$U_i^C \equiv \sum_{\ell \in \mathcal{L}_i} u_{i\ell}^C \qquad \text{and} \qquad U_i^P \equiv \sum_{\ell \in \mathcal{L}_i} u_{i\ell}^P. \tag{2}$$

Note that $U_i^C - U_i^P = \sum_{\ell \in \mathcal{L}_i} s_{i\ell}$.

4.1 A One-Shot Deviation Principle

Because the action sets are finite and monitoring is public within layers, the one-shot deviation principle applies to verify sequential rationality.[1] Thus it suffices to check that no player gains from deviating in a single period and then returning to the prescribed strategy.

4.2 Middle-Power Incentive Compatibility with Full Spillover

We start with the benchmark $\alpha = 1$ (perfect reputational portability).

Proposition 1 (SPE incentive constraint under full spillover $\alpha = 1$)**.** *Suppose Assumptions 1 and 2 hold with $\alpha = 1$. Fix a history with all reputations good. For any middle power i and any layer $\ell \in \mathcal{L}_i$, consider a one-shot deviation in which i plays D in ℓ (while others play C) and accepts a transfer $b_{i\ell}^t$. Under multiplex grim trigger, i does not deviate iff*

$$\pi_{i\ell} + b_{i\ell}^t \;\le\; \frac{\delta}{1-\delta}\left(U_i^C - U_i^P\right) \;=\; \frac{\delta}{1-\delta} \sum_{\ell' \in \mathcal{L}_i} s_{i\ell'}. \tag{3}$$

If (3) holds for all i and all $\ell \in \mathcal{L}_i$ at all histories with all good reputations, then the multiplex grim trigger profile is sequentially rational for middle powers.

[1] See standard references such as Fudenberg and Tirole (1991).

Proof. Fix i and $\ell \in \mathcal{L}_i$ at a history with all reputations good. If i follows the strategy, it cooperates in every layer forever, yielding present value

$$V_i^{\text{coop}} = \sum_{t=0}^{\infty} \delta^t U_i^C = \frac{U_i^C}{1-\delta}.$$

If i deviates only in layer ℓ at the current period, its current-period payoff increases by $\pi_{i\ell} + b_{i\ell}^t$ relative to cooperation (since the cooperative payoff U_i^C already includes $u_{i\ell}^C$). Under $\alpha = 1$, the deviation sets $r_i^{t+1} = B$ and triggers punishment in *all* layers in $\mathcal{L}_i$ from next period onward, yielding per-period payoff U_i^P thereafter. Thus the deviation value is

$$V_i^{\text{dev}} = U_i^C + (\pi_{i\ell} + b_{i\ell}^t) + \sum_{t=1}^{\infty} \delta^t U_i^P = U_i^C + (\pi_{i\ell} + b_{i\ell}^t) + \frac{\delta}{1-\delta} U_i^P.$$

The no-deviation condition $V_i^{\text{coop}} \geq V_i^{\text{dev}}$ is equivalent to

$$\frac{U_i^C}{1-\delta} \geq U_i^C + (\pi_{i\ell} + b_{i\ell}^t) + \frac{\delta}{1-\delta} U_i^P,$$

which rearranges to (3). Sequential rationality for punishers follows from Assumption 2: excluding a bad-reputation player weakly improves punishers' payoffs relative to inclusion, so no punisher has a profitable one-shot deviation in the punishment phase. By the one-shot deviation principle, the strategy profile is sequentially rational for middle powers whenever (3) holds for all relevant (i, ℓ). □

4.3 Comparative Statics: Multiplex Degree and Excludability

The enforcement power comes from $U_i^C - U_i^P = \sum_{\ell' \in \mathcal{L}_i} s_{i\ell'}$. This immediately yields two rigorous comparative statics.

Corollary 1 (Degree amplification under full spillover)**.** *Suppose $\alpha = 1$ and for a given i there exists $\underline{s}_i > 0$ such that $s_{i\ell} \geq \underline{s}_i$ for all $\ell \in \mathcal{L}_i$. Then for any layer $\ell \in \mathcal{L}_i$,*

$$\pi_{i\ell} + b_{i\ell}^t \;\leq\; \frac{\delta}{1-\delta}\, d_i\, \underline{s}_i \quad \Rightarrow \quad i \text{ rejects the deviation in layer } \ell. \tag{4}$$

Thus higher multiplex degree d_i raises the maximum bilateral inducement that can be resisted.

Proof. By Proposition 1, it suffices that

$$\pi_{i\ell} + b_{i\ell}^t \leq \frac{\delta}{1-\delta} \sum_{\ell' \in \mathcal{L}_i} s_{i\ell'}.$$

Under the hypothesis $s_{i\ell'} \geq \underline{s}_i$ for all $\ell' \in \mathcal{L}_i$, we have $\sum_{\ell' \in \mathcal{L}_i} s_{i\ell'} \geq d_i \underline{s}_i$, yielding (4). □

Corollary 2 (Excludability is the enforcement primitive)**.** *Under $\alpha = 1$, increasing $s_{i\ell}$ in any layer (i.e. making the club good in that layer more excludable/valuable, or making exclusion more painful) weakly relaxes the incentive constraint* (3) *for* all *deviations by i in* any *layer.*

Proof. Immediate from (3): the right-hand side depends on $\sum_{\ell' \in \mathcal{L}_i} s_{i\ell'}$, so increasing any term increases the continuation loss from reputational collapse. □

4.4 Partial Spillover: $\alpha \in [0, 1)$

We now prove the more general case in which cross-issue punishment is imperfect.

Proposition 2 (Incentive constraint under partial spillover)**.** *Suppose Assumptions 1 and 2 hold with spillover probability $\alpha \in [0, 1]$. Consider a one-shot deviation by i in a specific layer $\ell \in \mathcal{L}_i$. Then i does not deviate iff*

$$\pi_{i\ell} + b^t_{i\ell} \;\leq\; \frac{\delta}{1-\delta}\left(s_{i\ell} + \alpha \sum_{\ell' \in \mathcal{L}_i \setminus \{\ell\}} s_{i\ell'}\right). \tag{5}$$

Proof. If i cooperates forever, its value is $U^C_i/(1-\delta)$ as above. If i deviates only in layer ℓ at the current period, it gains $\pi_{i\ell} + b^t_{i\ell}$ immediately. From next period onward, the deviated layer ℓ punishes with certainty, while each other layer $\ell' \in \mathcal{L}_i \setminus \{\ell\}$ punishes independently with probability α. Hence the expected per-period payoff from $t+1$ onward equals

$$\mathbb{E}[U^{\text{cont}}_i] = U^C_i - s_{i\ell} - \alpha \sum_{\ell' \in \mathcal{L}_i \setminus \{\ell\}} s_{i\ell'}.$$

(Each punished layer ℓ' reduces i's payoff by $s_{i\ell'} = u^C_{i\ell'} - u^P_{i\ell'}$ relative to the cooperative path.) Therefore the deviation value is

$$V^{\text{dev}}_i = U^C_i + (\pi_{i\ell} + b^t_{i\ell}) + \frac{\delta}{1-\delta}\left(U^C_i - s_{i\ell} - \alpha \sum_{\ell' \in \mathcal{L}_i \setminus \{\ell\}} s_{i\ell'}\right).$$

Imposing $V^{\text{coop}}_i \geq V^{\text{dev}}_i$ and simplifying yields (5). As before, punishers have no profitable one-shot deviation by Assumption 2. □

Remark 1. *Equation* (5) *is the cleanest formal expression of the "connective tissue" claim: increasing α (better information-sharing / willingness to condition across issues) and increasing overlap $\sum_{\ell' \neq \ell} s_{i\ell'}$ both expand the bribe-resistance region.*

5 The Hegemon's Bribe Problem and the Divide-and-Rule Tax

To connect the middle-power enforcement result to hegemonic strategy, we characterize the *minimum* transfer needed to induce deviation, then embed it in a simple hegemonic optimization.

5.1 Minimum Bribe to Induce a Deviation

Define the *multiplex continuation loss* from deviating in layer ℓ:

$$\Phi_{i\ell}(\alpha) \equiv s_{i\ell} + \alpha \sum_{\ell' \in \mathcal{L}_i \setminus \{\ell\}} s_{i\ell'}. \tag{6}$$

Proposition 2 states that i resists deviation iff $\pi_{i\ell} + b^t_{i\ell} \leq \frac{\delta}{1-\delta}\Phi_{i\ell}(\alpha)$.

Lemma 1 (Minimum inducement)**.** *Fix (i, ℓ) and α. If the hegemon offers a transfer contingent on deviation in layer ℓ, then the minimum bribe that induces i to deviate (under multiplex grim trigger) is*

$$b^{\min}_{i\ell}(\alpha) = \max\left\{0,\ \frac{\delta}{1-\delta}\Phi_{i\ell}(\alpha) - \pi_{i\ell}\right\}. \tag{7}$$

Proof. By Proposition 2, i deviates iff $\pi_{i\ell} + b^t_{i\ell} > \frac{\delta}{1-\delta}\Phi_{i\ell}(\alpha)$. Solving for the smallest $b^t_{i\ell}$ satisfying this inequality yields (7) (allowing $b \geq 0$). □

5.2 A Simple Hegemonic Objective

Let $v_{i\ell} \geq 0$ denote the hegemon's per-period gross value from inducing i to bilateralize/defect in layer ℓ (improved terms, leverage, discriminatory access, etc.). If the hegemon offers $b_{i\ell}$ and i deviates, the hegemon's net value is $v_{i\ell} - b_{i\ell}$. If i does not deviate, the bribe is not paid.

A myopic but sharp characterization (sufficient for comparative statics) is:

Proposition 3 (Hegemon bribes iff profitable at the threshold)**.** *Under contingent transfers, the hegemon has a (weakly) optimal strategy of offering $b_{i\ell} = b^{\min}_{i\ell}(\alpha)$ if $v_{i\ell} \geq b^{\min}_{i\ell}(\alpha)$ and offering $b_{i\ell} = 0$ otherwise.*

Proof. If $v_{i\ell} < b^{\min}_{i\ell}(\alpha)$, any bribe that induces deviation yields negative net value and any bribe below $b^{\min}_{i\ell}(\alpha)$ yields no deviation and no benefit, so 0 is optimal. If $v_{i\ell} \geq b^{\min}_{i\ell}(\alpha)$, any inducing bribe $b \geq b^{\min}_{i\ell}(\alpha)$ yields net $v_{i\ell} - b$, maximized at the smallest such b. □

Remark 2. *Equations* (6)–(7) *formalize the "divide-and-rule tax": multiplex overlap and spillover (d_i, α) raise $b^{\min}_{i\ell}$ and thus shrink the set of (i, ℓ) for which bribery is profitable.*

6 Breaking a Club: Threshold Technologies and Tipping

Many coalitions are only effective above a minimum viable size. We capture this with a threshold m_ℓ.

Assumption 3 (Minimum viable cooperation in layer ℓ)**.** *In layer ℓ, if at least m_ℓ members remain in good standing and cooperate, then cooperating members receive $u^C_{i\ell}$ (club is "active"). If fewer than m_ℓ do, the club is "inactive" and all members receive a fallback payoff $\bar{u}_{i\ell}$ with $\bar{u}_{i\ell} \leq u^P_{i\ell} < u^C_{i\ell}$.*

Now consider a hegemonic "attack" on layer ℓ in some period in which all reputations are good. The hegemon can attempt to induce defections of a set $T \subseteq S_\ell$ simultaneously. If $|S_\ell \setminus T| < m_\ell$, then the club becomes inactive in that period (and, under grim trigger, remains so if members revert to fallback play).[2]

Define the number of defections needed to force inactivity:

$$d_\ell \equiv k_\ell - (m_\ell - 1).$$

Any T with $|T| \geq d_\ell$ collapses the club in that period.

Proposition 4 (Minimum cost to force a threshold collapse in a layer)**.** *Fix a layer ℓ and spillover α. Suppose the hegemon can offer contingent transfers to any subset of members in S_ℓ in a given period. To force the club inactive in that period, the minimum total transfer is*

$$\min_{T \subseteq S_\ell:\ |T| \geq d_\ell} \sum_{i \in T} b_{i\ell}^{\min}(\alpha) \;=\; \sum_{r=1}^{d_\ell} b_{(r)\ell}^{\min}(\alpha), \tag{8}$$

where $b_{(1)\ell}^{\min}(\alpha) \leq b_{(2)\ell}^{\min}(\alpha) \leq \cdots \leq b_{(k_\ell)\ell}^{\min}(\alpha)$ is the increasing order statistic of $\{b_{i\ell}^{\min}(\alpha) : i \in S_\ell\}$.

Proof. Because transfers are additive and there are no complementarities in inducing a given member's deviation beyond the threshold requirement, the hegemon's problem is a knapsack with unit "weights" (each induced defection counts as 1) and costs $b_{i\ell}^{\min}(\alpha)$. The minimum cost to obtain at least d_ℓ defections is achieved by selecting the d_ℓ cheapest members. Formally, suppose T is any feasible set of size d_ℓ. Replacing any $i \in T$ with some $j \in S_\ell \setminus T$ such that $b_{j\ell}^{\min} \leq b_{i\ell}^{\min}$ weakly reduces total cost while preserving feasibility. Iterating yields the set of the d_ℓ smallest $b_{i\ell}^{\min}$, proving (8). □

6.1 Network Structure and Bribery Targets

Combine Corollary 1 with Lemma 1: when surpluses are roughly comparable across layers, higher degree raises the minimum bribe. This pins down a sharp network implication: the hegemon optimally targets *peripheral* members in a multiplex.

Corollary 3 (Peripheral members are cheapest to pick off (under degree monotonicity))**.** *Suppose $\alpha > 0$ and $\Phi_{i\ell}(\alpha)$ is weakly increasing in d_i for members of S_ℓ (e.g. when $s_{i\ell'}$ are bounded away from zero and not negatively correlated with d_i). Then $b_{i\ell}^{\min}(\alpha)$ is weakly increasing in d_i, and Proposition 4 implies the hegemon's minimum-cost set T that collapses layer ℓ consists of the lowest-degree members of S_ℓ (ties broken by $\pi_{i\ell}$).*

Proof. If $\Phi_{i\ell}(\alpha)$ is weakly increasing in d_i, then by (7) so is $b_{i\ell}^{\min}(\alpha)$ (holding $\pi_{i\ell}$ fixed). Proposition 4 then implies the hegemon selects members with smallest $b_{i\ell}^{\min}$, i.e. lowest d_i (up to ties and heterogeneity in $\pi_{i\ell}$). □

[2] One can endogenize whether the club can be rebuilt; for present purposes the relevant object is the *minimum cost to force inactivity in the short run*, which is what targeted bribery exploits.

7 Interpretation: What Is (and Is Not) "Embedded Liberalism" Here?

The enforcement mechanism relies on three objects that can sound "liberal" in the IR vernacular:

1. **Portable reputation** (cross-issue conditioning),
2. **Information-sharing** (high α),
3. **Rule-like punishment** (automatic exclusion upon defection).

In this model these are not moral commitments; they are *enforcement technologies*. The liberal-looking machinery enters only as a solution to a realist incentive problem: in a world of hegemonic bilateral inducements, the only robust way for middle powers to sustain cooperation is to raise the continuation losses from deviation by making defections costly across a portfolio of interactions.

Crucially, Assumption 2 identifies the hard material core of the mechanism: without excludable club goods (or otherwise credible punishments), reputational conditioning is cheap talk.

8 Design Implications

The comparative statics are immediate and rigorous from (5) and (7):

- Raising α (better cross-club information and willingness to condition) increases $\Phi_{i\ell}(\alpha)$ and thus raises $b_{i\ell}^{\min}$.
- Increasing overlap (raising d_i or the size of $\sum_{\ell' \neq \ell} s_{i\ell'}$) increases bribe resistance.
- Increasing excludable surplus $s_{i\ell}$ (valuable, denyable benefits) is the most direct way to strengthen enforcement.
- Threshold clubs (m_ℓ) create tipping: once d_ℓ is large enough and/or the lower tail of $\{b_{i\ell}^{\min}\}$ is high enough, the hegemon's cost (8) increases sharply.

9 Conclusion

A one-shot perspective makes variable geometry look fragile. The multiplex repeated game shows why overlap matters: cross-issue punishment transforms many small cooperation problems into a single reputational enforcement system. The central object is the multiplex continuation loss $\Phi_{i\ell}(\alpha)$, which raises the minimum bribe required for deviation and thus imposes a divide-and-rule tax on the hegemon. The model also clarifies failure modes: low patience (δ), weak spillover (α), and non-excludable goods (small $s_{i\ell}$) collapse the mechanism back into bilateralism.

References

Abreu, D. (1988). On the theory of infinitely repeated games with discounting. *Econometrica*, 56(2), 383–396.

Fudenberg, D. and Tirole, J. (1991). *Game Theory*. MIT Press.

Keohane, R. (1984). *After Hegemony: Cooperation and Discord in the World Political Economy*. Princeton University Press.

Oye, K. (Ed.). (1986). *Cooperation Under Anarchy*. Princeton University Press.